

MULTI-IMAGE TIE-POINT DETECTION APPLIED TO MULTI-ANGLE IMAGERY FROM MISR

Jia Zong

Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Dr. Pasadena, CA, 91109, USA
Jia.Zong@jpl.nasa.gov

KEY WORDS: Tie-point, multi-image registration, relational-based feature matching, area-based matching

ABSTRACT:

An automatic tie-point (TP) detection algorithm for multi-image triangulation and registration is described. The algorithm uses a hierarchical approach that leads from the initial extraction of local image patches to the final TP detection on multiple imagery. The distribution of the TP detection is automatically adjusted to meet the needs of the triangulation or registration. Candidate point features are extracted based on the information of local image space only. A relational-based matching scheme using consistent labelling is designed for matching conjugate point features over multiple imagery. The final TPs are refined using the traditional area-based matching.

The algorithm was applied successfully to the in-flight georectification of global imagery from the nine pushbroom cameras of the Multi-angle Imaging SpectroRadiometer (MISR) instrument. TP match is accurate at 0.2 pixels and is closely successful on all cameras regardless a diverse range of geometric and radiometric distortions. The software was initially tested with simulated MISR data resampled from Landsat imagery and later applied to the production operation of the MISR in-flight georectification. The software was also applied to air-born MISR imagery. The results indicated the system could be adaptive to and effective on TP matching of space-born as well small scale air-born imagery.

1. INTRODUCTION

The automation of the traditional photogrammetric process that works with blocks of aerial imagery from metric frame cameras has been well researched [Tang and Heipke, 1996] and some commercialized [Walker, 1996]. In contrast with traditional processes, precise image registration of modern sensors with unique geometric and radiometric characters usually can not be performed on a standard softcopy station. Triangulation of special sensors is usually modeled according to sensors' specific geometry which consequently imposes unique requirements on the control point selection. With the ever increasing number of remote sensing satellites and air-born observation systems, automatic image triangulation and registration are becoming fundamental to remote sensing and photogrammetry. Automatic TP detection is an essential step to the success of image registration, especially large imagery.

MISR is one of the instruments aboard the Earth Observing System (EOS) Terra satellite [Diner, et.al. 1998]. Its purpose is to study the ecology and climate of the earth through the acquisition of systematic, global multi-angle imagery in reflecting sunlight. The instrument consists of nine pushbroom cameras pointing at discrete view angles (one nadir, four forward and four aftward, at angles of 0° , $\pm 26.1^\circ$, $\pm 45.6^\circ$, $\pm 60.0^\circ$, $\pm 70.5^\circ$ with respect to the surface, designated as An for nadir, Af, Bf, Cf, Df for forward set and Aa, Ba, Ca, Da for aftward set of cameras). Each camera has four spectral bands centered at 443, 555, 670, and 865 nm. Each channel consists of 1504 photo active pixels. Terra flies on a 705km sun-synchronous orbit, completes 233 unique paths every

16 days. During the level one standard processing of raw data, all 36 channels of imagery are required to be radiometrically calibrated, and georegistered to the pixel level (275m) with 95% confidence. The georegistered multi-angle and multi-band imagery are resampled onto a Space-Oblique Mercator (SOM) map projection and used in the level two geophysical parameter retrievals.

The standard georegistration processing routinely matches MISR imagery with a set of reference orbiting imagery (ROI) to correct for errors in the navigation data. ROI are non-cloudy MISR imagery resampled onto nominal Terra orbits. For continuing processing, ROI are created for all 233 orbits and nine cameras during the in-flight geometric calibration through an automatic triangulation and registration process. The details are described in [Jovanovic et. al., 1998]. In brief, a simultaneous bundle adjustment (SBA) with a spline model is used to triangulate the pushbroom cameras and correct for errors in navigation pointing. SBA adaptively sets knots along the orbit based on the navigation error type and the TP distribution. SBA also imposes requirements on the TP detection. First, TPs need to be consistently well distributed regardless surface type when applied globally. Second, the algorithm shall be well behaved with respect to geometric and radiometric distortions caused by diverse camera view directions. And last, the automatic detection needs to be reliable to ensure the triangulation accuracy.

This paper presents the TP detection algorithm used for MISR image triangulation and registration. Since automatic detection at the operation level imposes more complexity than investigations in the subject before [Agouris and Schenk, 1996], a hierarchical approach is designed to start with an initial match that isolates sen-

sor specifications from the generic TP detector. The objective is to formulate the TP detection requirements in the initial match and then perform multi-image matching according to the formulated descriptions that are independent of the object domain. Section two illustrates how initial match works. Section three describes the generic TP detector. The experiments are given in section four. The conclusion is provided in section five.

2. INITIAL MATCH

Initial match serves as the engine to control the distribution of TP detection and the inputs to the subsequent generic TP detector. It is entirely context driven, depending on the number and format of the input image, and the triangulation requirements. In practice, initial match creates TP detection candidates (TPC), on which TPs are detected by the generic TP detector. A TPC contains the following components: (a) the TPC location where TPs are desired; (b) multiple local image patches that contain overlap image; and (c) a set of constraints with respect to the TPC.

2.1. TPC Locations

A TPC location is a local region where one or a cluster of TPs to be detected. For MISR, TP detection needs to be sufficient to fulfill their task in the SBA, and yet efficient to balance for the computation expense. Without full knowledge on navigation performance pre-launch, TPC selection was designed on configurable grid evenly distributed along the orbit. In fact, TPC grid are defined on the standard Terra ground paths, where surface elevations and types (land or water) are provided by the MISR ancillary geographical product (AGP) [Lewicki & Zong, 1999]. In practice, special TPC locations within a grid cell are selected first, these are regions with rich surface features, either coastlines or hilly terrains, according to AGP. Figure 1 illustrates a segment of MISR ground path, where thick lines are the swath sides, dashed lines are




Figure 1: Tie-point candidate locations.

TP distribution grid, and little boxes are TPC locations with shaded boxes denoting to special TPC locations. Once a TPC is selected, local image patches and constraints are prepared for the generic TP detector to detect TPs. If enough TPs are detected in a grid, initial match will move on to the next grid. Otherwise, it will continue to select TPC within the current grid until either success or exhaustively failed.

2.2. Local Image Patches

Once a TPC location is defined, local image patches are extracted according to the initial relationship between image and surface. This is an easy task as most modern remote sensors are equipped with either satellite ephemeris or global positioning system along with inertial navigation systems (GPS/INS), which provides sufficient approximation for the initial match. In the case of MISR, an

image point intersection algorithm that employs collinary equation, as described in [Jovanovic, 1998], is used to project from the center of a TPC up to image space for all cameras. Local image centered around the intersections are extracted. Due to uncertainties in the navigation data and surface relief, local image patches must be large enough to ensure adequate overlap. For example, the dynamic error in the navigation data during a time interval for all MISR cameras to see the same ground feature could cause up to 10 pixel miss-coregistration from a pre-launch estimate. If generic TP detector requires 30 pixels multi-image overlap, local image patches must be over 50 pixels in one-dimension.

2.3. TPC Constraints

Given a set of local image patches, generic TP detection requires certain a prior knowledge to guide the detection. These constraints are grouped into two sets. The first set contains general information about the matching: (a) the number of matching patches, (b) minimum number of matched image patch per TP, and (c) the number of desired or *cluster* of TPs per image patch. In the case of MISR, there are total nine image patches to cover all nine cameras. A TP must be precisely matched on a minimum of five cameras in order to provide sufficient bundle constraint. At both sides of the orbital swath (cross-track direction), a cluster of two TPs per TPC were considered for more reliable attitude correction, whereas in the middle of the swath, one TP per TPC would be sufficient.

The second set of constraints regard to the intra-relation among image patches, required by the relational-based feature matcher. These include the pairing among the multi-image patches (generic TP detector starts with pair-wise matching), and additional constraints: (a) the scale factors of the imagery, (b) the accuracy of initial match, (c) the approximate surface relief, and (d) the camera pointing. These parameters need not to be accurate and can easily be derived from the general knowledge regard the sensor and the surface. Both sets of constraints are put together in an object called TP detection constraints and passed into generic TP detector.

3. GENERIC TP DETECTOR

Generic TP detector operates on multiple local image patches, governed by the TP detection constraints. First, interest points are extracted independently on every local image patch. Next, interest points are matched on pairs of local image patches and merged as potential TPs connecting all image patches. The TPs matched on the feature-base are then modified by area-based matchers and refined to sub-pixel locations. The process hierarchically employs the merits of feature-based and area-based matching, with the former providing efficient and reliable ties across the multiple imagery and the later enhancing the accuracy of the TPs.

3.1. Feature Detection

For multi-image matching, the interest operator needs to detect distinct point features that are invariant with respect to diverse geometric and radiometric distortions across the multiple image, regardless surface type.

After experiments with several operators, a variation of Forstner interest operator [Forstner, 1987] was used. In summary, Forstner interest operator detects meaningful point features such as corners

or local gravity center of image gray level. The original Forstner operator is computational expansive and requires an empirical threshold upon image conditions. We determined that Forstner operator should only be applied locally for detecting features that are distinct to local environment. Such detections are likely to be invariant relative to local image context even with large geometric and radiometric distortions. Selection of a 64 x 64 pixels square provides an adequate local region for interest point detection, which sets the dimension of the multiple local image patches for the generic TP detector. In addition to local application, a basic interest points are first selected at pixels with Robert gradients $\nabla g(i, j) > w_{mean}$, where w_{mean} is the mean magnitude of the Robert gradient over the local image patch. Forstner interest operator is then applied to basic interest points. This hierarchical approach not only reduces the computation significantly, but also improves the quality of detection with a simple weight threshold of one. Last, the window size for the suppression of local non-maximum interest values is set dynamically according to the number of the basic interest points as well as the diverse range of Forstner interest values within the local image patch. This dynamically controls the invariance of interest point detection, with respect to various global surface types that MISR covers, and various image qualities across MISR cameras.

3.2. Feature Match

A relational-based matching algorithm was designed that translates the problem of mapping of one set of features with another into a consistent labeling process.

3.2.1. Consistent Labeling Problem

According to [Haralick and Shapiro, 1993], an N -ary *consistent-labeling problem* (CPL) is a 4 tuple $CPL = (U, L, T, R)$. Component U is a set of M units $U = \{1, \dots, M\}$, which are the objects to be labeled. Component L is the set of possible labels. Component T is the *unit-constraint relations* over the unit set U . R is the *unit-label constraints* over the set $U \times L$ of unit-label pairs. A labeling of a subset $\hat{U} = \{u_1, u_2, \dots, u_N\}$ of U is a mapping $f: \hat{U} \rightarrow L$ from \hat{U} to L . A labeling f of a subset \hat{U} of units is consistent if whenever u_1, u_2, \dots, u_N are in \hat{U} and the N -tuple (u_1, u_2, \dots, u_N) is in T , then $[(u_1, f(u_1)), (u_2, f(u_2)), \dots, (u_N, f(u_N))]$ in R .

Now let A and B be two sets. Let $T \subseteq A^N$ be an N -ary relation over set A . Let $f: A \rightarrow B$ be a function that maps elements of set A into set B . The *composition* of T with f is defined by:

$$T \circ f = \{ (b_1, \dots, b_N) \in B^N \mid \text{there exists } (a_1, \dots, a_N) \in A^N \text{ with } f(a_i) = b_i, i = 1, \dots, N \} \quad (1)$$

Let $S \subseteq B^N$ be a second N -ary relation. A *relational homomorphism* from T to S is a mapping $f: A \rightarrow B$ that satisfies $T \circ f \subseteq S$. A relational homomorphism maps the elements of A to a subset of the elements of B having all the same interrelationships that the original elements of A had. The relational homomorphism problem fits the consistent-labeling model well. If we regard one data set A , such as a list of image features, as a unit set and another data set B , such as another list of features from a conjugate image patch, as the label set, the unit-constraint relation is simply the

relation T of the relational homomorphism problem with unit-label relation given by: $R = \{ [(u_1, l_1), (u_2, l_2), \dots, (u_N, l_N)] \mid (u_1, u_2, \dots, u_N) \in T \text{ and } (l_1, l_2, \dots, l_N) \in S \}$.

3.2.2. Tree Search

To solve for a consistent labeling problem, we look for the set of all consistent labeling $f: U \rightarrow L$ that satisfy the constraints specified by T and R . In the context of feature matching, there are M features from one image patch, N features from a conjugate patch, and $M \leq N$. Choosing M set of features as the unit set, and N set of features as the label set, the labeling of the unit set to the label set constructs a problem space which can be represented by a tree with a depth of M . Each node in the tree represents one labeling or pairing of a unit to a label. Each branch from the root to the leaf

represents one of total $\sum_{i=0}^M \prod_{j=0}^{M-1} (N-j)$ possible branches, only one of them is a consistent labeling representing the correct match.

The key of feature matching is to search for the consistent labeling and reject the inconsistent labelling efficiently. Due to image distortions, it is unlikely for features extracted from different images to have the same attributes and relations. Therefore feature matching may contain ambiguous solutions. For example, a unit set of 10 features is matching with a label set of 12 features. Tree search finds no branch with consistent labelling but two branches are partial consistent, one with three consistent labeled nodes and the other with eight. Obviously the branch that contains the maximum number of potential matches leads to the correct match.

In the discipline of artificial intelligence, one of the basic heuristic methods for pruning trees is to expand on a node according to its benefit function, defined as:

$$f(n) = g(n) + h(n) \quad (2)$$

where $g(n)$ is the benefit from the root of the tree to the current node and $h(n)$ is the estimated benefit from the current node to the final leaf. For image matching, $g(n)$ is defined to be the number of consistent labelling from the root to the current node, and $h(n)$ is the number of potential consistent labelling from the current node to the leaf if the current node is expanded. In this process, the future benefit is estimated according to a future-error-table (FTAB). Each element in the FTAB represents the error a pairing that deviates from the labeling constraints specified by T and R . When this error is large enough, a labeling is considered to be impossible. When the total number of consistent labelling from both past and future is too small, we stop expand along the current branch and backtrack for another node to expand.

3.2.3. Consistent Labeling Algorithm

The algorithm for solving the consistent labeling problem is based on the forward checking tree search by Haralick and Shapiro [1993], modified for expanding along a potential branch according to its total benefit.

Consistent_labeling($U_list, U_patch, L_list, L_patch, T, FTAB, R$)

Select the first unit in the U list
 Return success if no more unit is left in the U list
 For each potential label according to FTAB
 Add current (u, l) to R and return success if this is the last unit in the U list
 Call forward_check to update FTAB w.r.t. current labeling
 If total match from current to future larger than a threshold
 Call Consistent_label to search along the branch
 Add current (u, l) to R if search along branch is success
 End If
 End For
 If total match from current to future larger than a threshold
 Call Consistent_label to search along branch
 Add (u, no_label) to R if search along the branch is success
 End If
 Return fail (no more expansion for this unit and down, backtrack)

3.2.4. FTAB for Feature Matching

For consistent labeling of M units matching with N labels, FTAB is a $M \times N$ array with each element represents the error of labeling unit u with label l . For interest point matching, the size of image search window is used to initialize FTAB as a binary map to prune non-potential labels for all units. In the case of MISR, local image patches are defined according to image navigation and the average elevation of the local surface. The static and dynamic errors in the navigation data propagate to an uncertainty of $[l_0 \pm \Delta l, s_0 \pm \Delta s]$ in the local image extraction. MISR cameras also contain a diverse range of pointing in the along-track direction ($0.0^\circ, \pm 26.1^\circ, \pm 45.6^\circ, \pm 60.0^\circ, \pm 70.5^\circ$), which creates image disparities for surface features that are off the local average elevation. Combining both factors, a search window is determined for each pair of local image patches and used to define candidate labels of all unit features.

Next, the error value of each candidate pair in the initial FTAB is replaced with the unary constraints. First, the local distinctness of an interest point is represented by its interest value w . Conjugate interest points surrounded by similar image patterns tend to have similar interest values for locally detected interest points. The first unary constraint for matching is defined as the relative difference of the interest values:

$$W_{sim} = \frac{w_l - w_u}{\min(w_l, w_u)} \quad (3)$$

The next unary constraint is the radiometric similarity of local images, defined by a cheap area-based similarity measurement:

$$R_{sim} = \frac{\sum_{f \in abs} \left\{ |img_l(r, c) - \overline{img}_l| - |img_u(r, c) - \overline{img}_u| \right\}}{\sigma_l} \quad (4)$$

where w defines a 5×5 similarity window centered at the interest point, $img(r, c)$ is the image value at pixel (r, c) , \overline{img} is the mean image value within the similarity window, and σ_l is the sigma of image values within the label similarity window. The total unary

constraint is the summary of the distinctness similarity and radiometric similarity, normalized by a factor of two.

Before a tree search starts, the unit list is ordered such that units with less labeling candidates are listed first to prune out a portion of unnecessary search tree. During a tree search, FTAB is updated according to a topological binary relationship between interest points. Topological relationship is used because local rotational distortion is relatively small for small scale imagery. For each pair of interest points, the topological binary distances are $D_{x_{ij}} = s_x \cdot |x_i - x_j|$ and $D_{y_{ij}} = s_y \cdot |y_i - y_j|$, where i and j are interest point indices, x and y are interest point coordinates, s_x and s_y are pixel resolution of the image. The topological binary constraints are:

$$B_{x_{ij}} = \frac{\left| (D_{x_{ij}})_u - (D_{x_{ij}})_l \right|}{(D_{max})_x} \quad B_{y_{ij}} = \frac{\left| (D_{y_{ij}})_u - (D_{y_{ij}})_l \right|}{(D_{max})_y} \quad (5)$$

where $(D_{max})_x$ and $(D_{max})_y$ are relaxed pixels for this constraint. For small scale imagery, they are about two to three pixels.

3.2.5. Tie Point Merge:

The relational-based feature matching results in a list of matched interest points for each image pair. Combining all match lists together is simply a sorting process. For example, in the case of MISR, local image patches from nine cameras are paired by either adjacent or every other adjacent cameras to provide closer geometric and radiometric similarity. This results in total 15 image pairs: DfCf, CfBf, BfAf, AfAn, AnAa, AaBa, BaCa, CaDa, DfBf, CfAf, BfAn, AfAa, AnBa, AaCa, BaDa. Assume interest point pair (0, 4) is on the match list of image pair DfCf, (0, 2) is a point pair on the list of image pair DfBf, a new TP with ($tp_id = 0$, $Df = 0$, $Cf = 4$, $Bf = 2$) is created. TP merge also provides a reliability assurance. If (4, 1) is also a point pair on the match list of camera pair CfBf. The insertion of this pair would result in an inconsistency in the TP table, therefore any match related to TP $tp_id = 0$ is discharged as a blunder. The algorithm used to sort the pair-wise matched interest points and create TPs across multiple image patches is the determination of equivalent class from Numerical Recipes in C by Press. et. al [1992].

3.3. Precise Match

The accuracy of relational-based feature matching is about two pixels due to both image distortions and the relaxed geometric constraints. Area-based matchers then refine the accuracy to sub-pixels. A refined TP must be precisely matched on a minimum number of local image patches. MISR triangulation requires a TP be precisely matched on a minimum of five cameras with 0.2 pixels accuracy, while the rest of cameras could be relaxed. If a minimum number of TPs are precisely matched according to the cluster requirement from triangulation, precise match will stop to refine the rest of feature-matched TPs.

The main criteria for precise matching is high precision and reliability. First, cross-correlations are computed for each TP. The template window for a TP is centered at the interest point with the

highest interest value. To ensure reliability, the maximum correlation location may not be more than two pixels away from the initial feature on the search image. In addition, the correlation values of the eight direct neighborhood pixels around the maximum should also be locally large indicating the center pixel is on the hill of a correlation surface. The 3 x 3 local correlation hill is represented by a quadratic two-dimensional polynomial with the maximum of the polynomial as the final solution of the correlation. The σ for correlation is better than half a pixel. To further improve, least-square matcher (LSM) [Gruen and Baltsavias, 1987] is followed. The same template camera is used. LSM sequentially matches the template with every conjugate image patch instead of simultaneously applied to multiple image. This way, the matcher is separated from the engine that controls the matching legislation. The correction of LSM is only allowed within one pixel for high reliability. The σ for LSM is less than two tenth of a pixel. Though this combination of correlation and LSM algorithms creates an expensive area-based matcher, it still has a high performance due to the quality of the input TPs and the small search windows.

4. EXPERIMENTS

4.1. Pre-launch Tests

In prelaunch development, an orbit of simulated MISR imagery that runs through the north America continent was created, based on similar bands of Landsat Thematic Mapper (TM) imagery along with a registered DEM. The simulation process is described in [Lewicki et. al., 1994]. The TP detection algorithm was implemented and tested with the simulated image. The results were applied to SBA for evaluations against various orbit error models.

Due to the dynamic errors in the simulated orbit, SBA requires 20 well distributed TPs per 512 image lines. SBA also shows cluster of TPs at both sides of the swath are not necessary. TP detection was configured accordingly. Figure 2 illustrates the number of detected TPs per SOM block along the swath. A MISR ground path on SOM projection is reported on 180 blocks from north to south. Each SOM block is 140.8km long in the along-track direction and contains 512 lines of projected image. The land coverage of the simulated MISR imagery is from block 51 to 77. Note the TPs in the start and end blocks are less since they are partial blocks. The overall TPs could be increased if configured to. On average, there is over one TP detected per TPC per MISR camera regardless surface type, with performance slightly better on hilly terrain. The detection is slightly weakened on the most oblique D cameras, covering only 80% of the grids set by initial match.

Generic TP detector was intensively tested with the simulated MISR imagery that cover different surface types and are resampled from LandSat imagery of different times. It worked effectively and consistently if initial match would supply local image patches with over 30 pixels overlap in both dimensions. Though most parameters are configurable, the system performed well with the defaults in all tests. On average, there are 15~20 interest points detected on a local image patch, and among them 10~15 interest points are matched in feature base consistently (MISR D cameras are likely on the lower end). The experiments also show that the detection of invariant interest points over multi-image are critical

and the relational-based matcher is very reliable.

The accuracy of the matching was assessed in three ways. First, TP symbols were superimposed on imagery and visually examined by zooming into the pixel level for measurements. Second, "truth" TPs were simulated for all detected TPs, by projecting from one precisely matched camera down to the surface DEM, then backward up to the rest of cameras. The projection used the simulated navigation data with no added errors, in contrast with the "measured" navigation data used in TP detection and SBA. The mean differences between the detected and simulated TPs are within 0.2 pixels for LSM points, 0.4 pixels for correlation points, and about 2 pixels for feature TPs, all with a small sigma. Finally, SBA would evaluate the residuals of the TPs after trigulation and compare with the errors added in the orbit model. The results verified the uncertainties of various types of matching and the sufficiency of TP detection with respect to the need of triangulation.




Figure 2. Number of detected TPs per SOM block.

4.2. Post-launch Operation

TP detection was tested with MISR data after Terra launched in late 1999. With little change, it was set into production at an average rate of processing one day-side orbit per 45 minutes and run automatically upon data coming. Figure 3 is an example of TP detection on MISR orbit 2578 over south Greece for camera Df and Aa, where the red symbols are TPs matched with LSM, the yellow ones are matched with correlation, and the green ones are feature matched. Note that image Df has less match as the camera looks through a long path in the atmosphere and image becomes fuzzy. Quality assessment statistics, such as the plot shown in Figure 2 as well as other SBA feed backs, are created automatically for each orbit production. The overall production ended within a few months and successfully covered all Terra paths. Though there are a few paths with very few TP detections, the investigation showed they all due to the imagery is too cloudy according to the MISR cloud detection standard.

4.3. Adoption to AirMISR

The pushbroom AirMISR flies on 20km above the surface. Comparing with MISR, both its position and orientation are unstable, causing large image distortions. Generic TP detector is applied to imagery resampled on the ground using initial navigation. When initial navigation provides near 30 pixels overlaps, generic TP detector works fine as to MISR, regardless the much worst image distortions. When image pointing is worse, initial match would



Figure 3. TP detection on orbit 2578 for camera Df and Aa.

impose match in pyramid or an operator would need to measure one TP on all images to supply a correction to initial match.

5. CONCLUSION

A multi-image TP detection algorithm was designed and successfully applied to the triangulation of MISR imagery globally. The hierarchical approach leads the detection to a maximum coverage across all images and all surface lands. Generic TP detector is flexible in application, very reliable benefit from the relational-based matcher, and accurate with the enhancement from the area-based matcher. The system could be potentially applied to or adapted to different types of multi-image matching.

6. ACKNOWLEDGMENTS

The author gratefully acknowledges the efforts of the MISR Science Data System Team: Graham W. Bothwell, Mike A. Bull, Earl G. Hansen, Veljko M. Jovanovic, and Mike M. Smyth. This research is being carried out at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. For further information, see MISR web site at: <http://www-misr.jpl.nasa.gov>.

7. REFERENCES

Agouris, P., & T. Schenk, 1996, Automatic Aerotriangulation Using Multiple Image Multipoint Matching, *PE&RS*, 62(6):703-710.

Diner, D. J., J. C. Beckert, T. H. Reilly, T. P. Ackerman, C. J. Bruegge, J. E. Conel, R. Davies, S. A. W. Gerstl, H. R. Gordon, R. A. Kahn, J. V. Martonchik, J.-P. Muller, R. Myneni, B. Pinty, P. J.

Sellers, and M. M. Verstraete, 1998, Multi-angle Imaging Spectro-Radiometer (MISR): instrument description and experiment overview, *IEEE Trans. Geosci. Rem. Sens.*, 36(4), 1072-1087.

Forstner W. and E. Gulch, 1987, A fast operator for detection and precise location of distinct points, corners and centers of circular features, *ISPRS Intercom. Workshop*, Interlaken, 281-304.

Gruen, A. W. and E. P. Baltsavias, 1987, High-Precision Image Matching for Digital Terrain Model Generation. *PRS*, 42:97-112.

Haralick, R. M., and L. G. Shapiro, 1993, *Computer and Robot Vision*, Addison-Wesley Publishing Company, Inc., Vol. II, p379.

Jovanovic, V. M., M. M. Smyth, J. Zong, R. Ando and G. W. Bothwell, 1998, MISR Photogrammetric Data Reduction for Geophysical Retrievals, *IEEE Trans. Geosci. Rem. Sens.*, 36(4), 1290-1301.

Lewicki, S. A. and J. Zong, 1999, Level 1 Ancillary Geographic Product Algorithm Theoretical Basis, *JPL D-13400. Rev. B*.

Lewicki, S. A., M. M., Smyth, V. M. Jovanovic, and E. G. Hansen, 1994, A Simulation of EOS MISR Data and Geometric Processing for the Prototyping of the MISR Ground Data System, *IGARSS, vol III*, Pasadena.

Tang, L., and C. Heipke, 1996, Automatic Relative Orientation of Aerial Images, *PE&RS*, 62(1):47-55.

Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, 1992, *Numerical Recipes in C - The Art of Scientific Computing*, 2nd Edition, Cambridge University Press, p345.

Walker, S., 1996, Digital Photogrammetric Workstations 1992-1996, *Int. Arch. of Photog. & RS*, 31(B2):384-395.